

DIPP: Diffusion of Privacy Preferences in Online Social Networks

Albert Mwanjesa, Onuralp Ulusoy, and Pinar Yolum

Universiteit Utrecht, Utrecht, The Netherlands,
Princetonplein 5, Utrecht, 3584 CC Utrecht
a.mwanjesa@gmail.com, o.ulusoy@uu.nl, p.yolum@uu.nl

Abstract. Ensuring the privacy of users is a key component in various computer systems such as online social networks, where users often share content, possibly intended only for a certain audience but not others. The content users choose to share may affect others and may even conflict with their privacy preferences. More interestingly, individuals sharing behaviour can change over time, which indicates that privacy preferences of individuals affect others and spread throughout the network. We study the spreading of privacy preferences in online social networks with information diffusion models. Using multi-agent simulations, we show the dynamics of the spread and study the factors (e.g., trust) that influence the spread.

Keywords: privacy, diffusion, multi-agent systems

1 Introduction

With the rise of social media and the prevalence of smartphones, humans are more connected online than they have ever been before. In general, this means people can easily communicate with each other no matter their location, physical social circles and so on. The communication is many times done on public platforms, such as online social networks (OSNs) in the form of sharing content, such as pictures, videos, and so on. The content shared by users can reveal their private information, either explicitly or implicitly.

The problem of privacy preservation in OSNs can be viewed from various angles due to its interdisciplinary properties. Dupree *et al.* [4] investigate privacy personas in online social networks. Using a survey, data are gathered on user behaviour towards privacy and security. A cluster analysis shows that individuals can be categorized in one of the following categories: Fundamentalists, Lazy Experts, Technicians, Amateurs and the Marginally Concerned in terms of their knowledge in privacy and motivation in sharing content. Interestingly, a source of influence in privacy decision-making is an agent's trust towards their neighbours. A qualitative study by Lampinen *et al.* [8] have found that privacy management in content sharing decisions is mostly based on trust. Users expect each other to understand the way they want to represent themselves in OSNs. To ensure

this, users are said to utilize mental/behavioural strategies as well as preventive, corrective and collaborative strategies.

Researchers have investigated the problem of privacy preservation in OSNs using approaches based on multi-agent systems. Kokciyan and Yolum [7] use commitments between agents and norms [3] in a system to prevent and detect privacy violations. Ulusoy and Yolum in [12] have shown that online social networks can leverage social norms as they emerge from personal norms. Understanding the evolution and spread of privacy norms is important because they can serve as an important tool to understand and avoid privacy violations. Privacy can also be violated by other users as well when content is shared that reveals private information about someone without their consent. Sharing content that is owned by multiple users can lead to leaking of private information. Social norms can be used to guide privacy management, but personal privacy preferences are unclear to other users. It is evident that one's level of openness is not the same for all and violations can occur. It is, thus, useful to understand how privacy preferences spread through a social network.

To model the spread of privacy preferences, inspiration is taken from information diffusion research. Information diffusion research tries to understand how information spreads in, for example, OSNs[5]. The diffusion of information has often been modelled using epidemic models. Epidemic models can be seen as state machines governed by probabilities of transitioning from state to state. They contain three common states: (S)usceptible, (I)nfectious, (R)ecovered. There are different compositions given these states, from SI, SIS and SIR to many more variations. Originally, epidemic models were used to model the spread of infectious diseases, hence the name. But they were adopted for information diffusion, as it was theorized that information can be infectious [9]. These models are flexible, as illustrated in work by Cannarella *et al.* in [1]. The researchers used an epidemic model to predict adoption and abandonment of a social media platform. Here, we argue that privacy preferences can be derived from the privacy decisions users make. The context of content shared by users describes their privacy preference, as elaborated later. We study the following research questions.

- RQ1: How does a privacy preference of posting content that fits in a given context spread?
- RQ2: Who are the most influential figures in the spreading process of the privacy preferences?
- RQ3: What is the influence of trust modelling in the spread of privacy preferences?

In order to answer these research questions, we aim to contribute to the literature by (i) investigating whether epidemic models are accurate models to model the spread of privacy preferences (ii) analysing the influence of different factors in a social network on this diffusion process, (iii) simulating the diffusion of privacy decisions, and (iv) providing a method for users to protect themselves against opposed privacy preferences, namely using trust modelling.

The rest of this paper is organized as follows. Section 2 develops our model for privacy preference diffusion. Section 3 explains details for realizing the model

as an agent-based system. Section 4 depicts our experiments and results. Finally, Section 5 discusses the work with pointers to future work.

2 Modelling privacy preference diffusion

We model an online social network (OSN) as a multi-agent system, where each agent represents a user that shares content in a given context. In real life, when a user observes a piece of content on an OSN, she perceives different properties of the content that leads the user to react such as to appreciate by in the form of a *like*, to dislike it, or at times to ignore the content. Imagine a photo of someone taking a self-portrait photo of a user at the pet zoo with a llama in the morning. The description of the scene of the photo captures its *context*. This context defines location (i.e., zoo), people in the picture (i.e., user), what is happening in the picture (i.e., posing with a llama) and the time of day (i.e., the morning).

In this work, the context of the content shared by a user on an OSN is assumed to reveal the privacy preferences of that user on said OSN. This follows from the assumption that a user only shares content that they want to be seen by their friends on an OSN.

Context, in this study, is described using three locations and four times of day. The four times of day are *morning*, *afternoon*, *evening* and *night*. The three locations are at *work*, *the beach* and *the mall*. The contexts are represented as 2-tuples with 12 possible combinations, for example: $\langle \textit{night}, \textit{work} \rangle$.

Definition 1 (Privacy preference). *A context description that can be ascribed to the content shared by an agent with the assumption that agents only share content they do not find to be private.*

Privacy violations in OSNs occur when content reveals information about a user that the user would want to keep private. Opposed privacy preferences can be deemed to describe content that a user believes should be private. In this project, the focus is on the unopposed privacy preferences. However, to capture privacy violations, the spread of opposing privacy preferences is modelled in two different ways. It is assumed that opposing privacy preferences either spread the same way as their unopposed counterparts or are static. Privacy violations can occur when a user shares content that is co-owned by another user and, in doing so, reveals private information of said user. Using privacy violations, we can model trust between agents on an OSN. Using the model, we will investigate if the use of trust values can help agents protect themselves from opposed privacy preferences.

These are the underpinnings of this research; privacy preferences are infectious via the medium of the content sharing habits of friends. Users perceive friends sharing certain content and are deemed more likely to adapt the same habits. In line with information diffusion models, we believe that privacy preferences are, thus, infectious. This section outlines the global theories on how these infections come about and how they could be influenced.

Definition 2 (Co-owned content). *Content is co-owned when it represents more than one user. Representation could be by appearance in a picture, sound of a voice in a video or audio recording, explicit tagging of users and more.*

Definition 3 (Privacy violation). *A privacy violation in an OSN is an instance in which content is shared without the explicit consent of a user. Furthermore, this content is in a context that the user opposes sharing content of.*

The next sections present a model that simulates the spread of privacy preferences and the varying factors that are believed to affect this dynamic in OSNs. Opposed privacy preferences, co-ownership, privacy violations and trust are believed to impact the diffusion of privacy preferences in reality.

2.1 SIR model for privacy preference diffusion

The SIR model is an epidemic model that is governed by an infection rate and a recovery rate. An agent in the model can be in one of these three states: (S)usceptible, (I)nfected and (R)ecovered. An agent starts in a susceptible state, meaning it can be infected by an infectious entity. When the infectious entity infects an agent, the agent moves to the infected state. If and when an agent recovers from this infectious entity, the agent is in the recovered state. Being in the recovered state is synonymous with being immune to the infectious entity, as there is no state transition out of the recovered state. The state dynamics of the SIR model are governed by an infection rate i and recovery rate r , as visualized in Figure 1a.

We use the SIR model for privacy preference diffusion for OSNs where the semantics of infection are applied to privacy. Thus, (S)usceptible state means that agents can imitate certain behaviour if they witness it from their neighbours; e.g., observe content that depict certain privacy preferences. (I)nfected state denotes that the agent has observed another agent to whom it connected, imitated its behaviour by sharing a content with the said privacy preference. This will be reflected by an infection rate for that privacy preference. (R)esistant state denotes that the agents have recovered from an infection of a privacy preference and because they are now immune to the said privacy preference, they will not share content aligned with said privacy preference. An agent can recover from an infection when content they have shared has negative consequences on them, or when sharing certain types of content becomes less fashionable. An example to the first case is content shared by an agent can lead to privacy violation of a friend, having a negative impact on their relationship. For the second infection recovery case, one infamous example is the challenge trends on social media, where sharing a video of doing a specific task becomes viral for a time period but in time it becomes out of fashion and people recover from this infection. Thus, the states determine the likelihood of an agent to share content described by a certain context. These state transitions are governed, mainly, by the infected rate i and recovery rate r and can be seen in Figure 1a. From these models, the key measurements are population proportions for each state, as is common in

diffusion modelling. The epidemic model starts with multiple infectious entities. To make the model more realistic, in terms of modelling the spread of privacy preferences, concepts such as opposing infectious entities, trust and co-ownership of content are added. There could be different models that might have worked in this study, such as the DeGroot model [2], but considering the flexibility, dynamic trust values, and foundation in previous research, we opt for the SIR Model.

2.2 Trust

Research in privacy has shown that privacy management in content sharing of OSNs is largely dependent on trust among users [8]. Inspired by that, the DIPP model implements one-dimensional trust modelling. One-dimensional trust refers to an agent’s belief in another agent to complete a task with one measurement of success [10]. The task delegated in the DIPP model is for a neighbour to respect an agent’s privacy. Trust in the DIPP model is influenced by interactions between agents of which they are two: perceiving another agent’s shared content and being a co-owner of the content shared by another agent.

To model trust, the beta distribution is used [6]. This probability distribution is governed by two variables α, β that can be used to capture interaction outcomes [10]. We use α to capture the cases where a privacy violation has not taken place and β for cases where there is a privacy violation. This leads to two values influenced by the two interactions, but based on the same delegated task. We differentiate between the two types of privacy violations as follows. A *Level 1* privacy violation occurs when an agent perceives another agent sharing content that portrays a privacy preference that they oppose. A *Level 2* privacy violation occurs when an agent shares content without a co-owners’s consent. The trust values tr_1, tr_2 and severity levels s_1, s_2 represent the trust value for a type of privacy violation and the perceived severity of said violations respectively. Finally, the two values for the two types of violations are integrated into one with weights an agent uses to express what they see as more severe

$$tr = s_1 \times tr_1 + s_2 \times tr_2$$

with $s_1 + s_2 = 1$. The trust value directly influences the infection rate as can be seen in Figure 1b. In this project, the severity levels are static over all simulations.

2.3 Co-ownership

Users should be more likely to be infected with a privacy preference when they are co-owner of the content being shared. To this end, $1.5 > co > 1.1$ is a random number that represents how impressionable an agent is, with regard to sharing content they have doubts over. As with h , co is a random number to capture the fact that not all agents in a social network can be persuaded on the same level and this resilience is also not static overtime. The state dynamics given content co-ownership are visualized in Figure 1c.




(a) State transitions in the SIR model with infection rate i and recovery rate r (b) State transitions in the SIR model with the infection rate i and recovery rate r with tr representing the trust in the neighbour trying to infect the agent(c) State transitions in the SIR model with the infection rate i and recovery rate r with co in the case of an infection attempt via co-ownership

Fig. 1: State transitions in the DIPP model

3 Realization of the DIPP model

In the DIPP model, an agent represents a user on an online social network that shares content from which other agents can derive said agent's privacy preferences. To this end, each agent must have the ability to share content and all their friends should be able to perceive the content. A schematic view of an agent in the DIPP model can be seen in Figure 2a. We have implemented the DIPP model on top of the Mesa library for Python 3¹. The time steps are discrete.

Agent initialization: Each agent is assigned privacy preferences on the pro and anti sides. For each privacy context, the agent is randomly assigned either the state susceptible or infected, with a probability of 0.5. The resistant state is ignored here as it is a final state and would thus limit the dynamics in the experiments.

Definition 4 (Pro side). *The epidemic of privacy preferences that are supported by agents. These are the privacy preferences that underlie the infected agents' content sharing habits.*

Definition 5 (Anti side). *The epidemic of privacy preferences that are opposed by agents. Infected agents on this side believe that the content, described by the privacy preference they are infected with, should not be shared.*

¹ <https://mesa.readthedocs.io/en/master/overview.html>

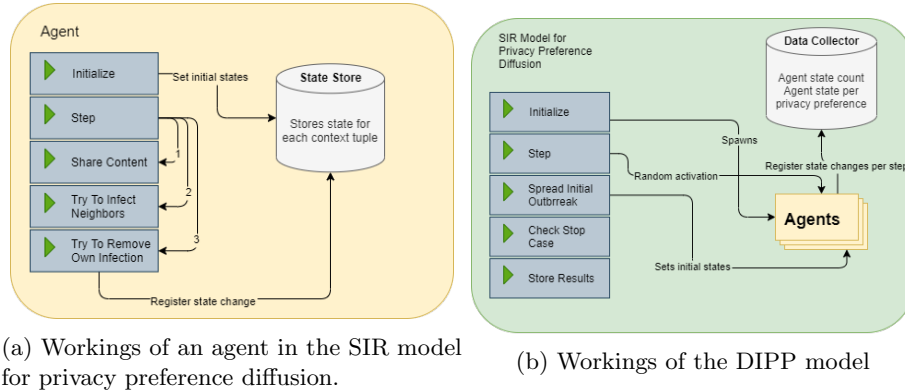


Fig. 2: Schematic view of the components in an agent and the DIPP model

For the epidemic dynamics to work, infection rates and recovery rates have to also be assigned for each of the privacy preferences. These can be single numeric values that hold for either pro-side or anti-side. However, the rates can also be defined as a mapping between a context and a numeric value, allowing for precise custom assignment of rates for any experimental setting.

Agent Actions: Every time an agent takes a *step* in this model, they do at least three things: share content if possible on pro-side, share content if possible of anti side and try to recover from an infection. An agent shares content by randomly choosing one of the contexts they are infected with on the pro-side. Once content is shared, neighbours can perceive and become infected by the context of the content. The infection rate is dependent on opposing privacy preferences, trust and co-ownership. The final part of a step is the *attempt* to recover from an infection. In this part, a context tuple is chosen randomly from all the context tuples for which the agent is infected. The context tuple’s recovery rate is the probability that the agent will recover in this step.

Simulations: The DIPP model simulates the spread of privacy preferences on an OSN. A schematic view of the DIPP model can be seen in Figure 2b. For a researcher, the model is the interface to interact with to run any simulations. It creates agents for each node on the social network provided to it. The model collects for each context tuple the number of agents in each of the three states. It also collects data on the number of state changes per step. Finally, the model also keeps track of the infection chains for each context tuple. The model also provides an interface to customize the initial outbreak of privacy preferences. Furthermore, it also provides the ability to specify the state of a specific agent regarding a specific context tuple. It is possible to run a simulation for any number of steps. However, there is also an option to run a simulation until the state dynamics have reached a stable condition. Stability is reached when a predefined number of steps has passed and there have been fewer state changes than a predefined fraction of the population.

4 Evaluation

We use simulations to evaluate if our proposed DIPP model can indeed model the spread of privacy preferences through privacy decisions. Given the data from simulations with different settings, changes in the state dynamics and agent influence on the diffusion processes can be evaluated.

4.1 Measurements

In information diffusion modelling, influence has always been important. Sometimes influential users in a network are sought after for different reasons. It might be because they could be used to promote information on the network. In other cases, they might be deemed *bad influences*. The *influence rating* of an agent is the count of each descendant proportional to the length of the shortest path between the agent and the descendant in the infection chain.

One of the goals of this project is to explore whether trust values can be used by agents to protect themselves against opposed privacy preference. To this end, the epidemic endings are measured. An *epidemic ending* of a privacy preference is the step at which no agent on the network is infected with said privacy preference. Similarly, an important goal is to explore the effect of trust in slowing down infections. To this end, the epidemic peaks are measured. An *epidemic peak* of a privacy preference is the maximum number of infected agents over the period of epidemic, infected with that privacy preference.


4.2 Experiments

For the purposes of answering the research questions of this project, simulations are run with the model described in the previous section using a real online social network data set called *Copenhagen Networks Study interaction data set*. This data set represents a multi-layer temporal network collected at the Technical University of Denmark from freshmen students at the same university [11]. Here, we make use of the Facebook friendship network, which is an undirected graph with edges representing friendships that last during the whole experiment in which the data were collected. The degree distribution of this network shows some similarity with the power law distribution, which is often regarded to be representative of social networks.

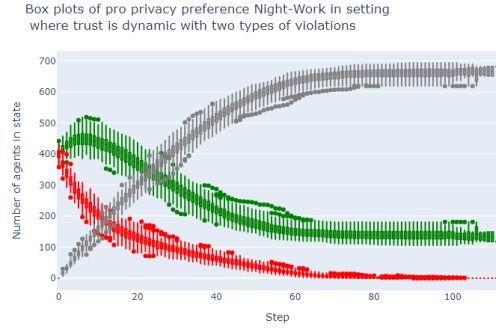
Some settings or setups are constant across all simulations. These settings will be explained first. After this, the specific simulations are put forth along with the expectations from these experiments.

Basic settings All the simulations use flat infection rates and recovery rates. This means that all the privacy preferences have the same rates on both pro and opposing sides. Three rates are chosen, $\{0.25, 0.50, 0.75\}$. These three values are varied over the four variables: pro-infection rate, pro-recovery rate, anti-infection rate and anti-recovery rate. These lead to a total of, at most, $3^4 = 81$ combinations of rates to run. In the case of disregarding the opposing preference dynamic, there is no need to include the anti infection and recovery rates.


Consequently, there are $3^2 = 9$ combinations of rates to run. Every simulation is executed 100 times to account for the stochastic nature of the model's dynamic. 100 repetitions is deemed to be the right compromise between correctness of the results and the run time of the experiments. This means that one experimental setting amounts to 900 or 8100 simulation runs. A simulation is run until a stable condition is reached. Stability is defined using a look back of 20 steps and a fraction of the population of 0.05. This means that a simulation ends when in the last 20 steps there have been fewer state changes than $0.05 \times 800 = 40$.



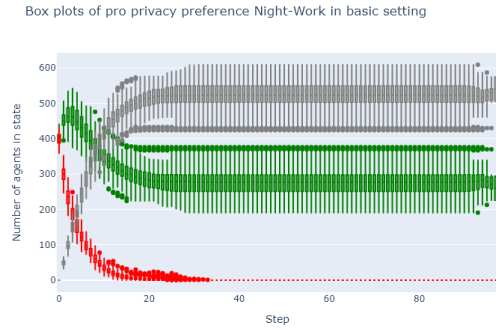
(a) pro-infection rate 0.25, pro-recovery rate 0.25, anti-infection rate 0.25, anti-recovery rate 0.25



(b) pro infection rate 0.25, pro-recovery rate 0.25, anti infection rate 0.25, anti recovery rate 0.25



(c) pro infection rate 0.75, pro-recovery rate 0.25, anti infection rate 0.25, anti recovery rate 0.75



(d) pro infection rate 0.25, pro-recovery rate 0.75, anti infection rate 0.75, anti recovery rate 0.25

Fig. 3: State dynamics of baseline simulations of the DIPP model. (S)usceptible (I)nfectious (R)ecovered

These basic settings are used to generate baseline data. The baseline experiment to show that:

- **1a** a privacy preference epidemic will last longer when the infection rate of said privacy preference is higher than its recovery rate,
- **1b** this setting will show a positive correlation between the degree of an agent on the network and its influence on the spread of privacy preference.

These assertions need to hold for the model to, at least, be a valid epidemic model. Thus, they are the foundations of the DIPP model. Furthermore, a node’s degree has been shown previously to correlate with the influence of that node in information diffusion [9]. That feature is expected to stay the same in the context of privacy preference diffusion as modelled here. In figures 3a, 3c and 3d, we can see aggregated state dynamics from the baseline simulation for the pro-side epidemic of privacy preference $< Night, Work >$. When all rates are equal, there is a steep increase in recovered agents and almost linear decrease in the number of susceptible and infected agents. When the anti-side epidemic is weak, Figure 3c, the number of infected agents increases after the start steeply, while the number of susceptible agents decreases steeply. It should be noted that when the anti-side epidemic is weak there comes a point in the simulation at which there are no more susceptible or infected agents. This is in contrast to when the pro-side epidemic is weak, Figure 3d. The roles in decreasing and increasing are reversed, but it is also noticeable that the final number of susceptible agents is higher and never goes to zero.

Simulation results indicate that a privacy preference epidemic will last longer when the infection rate of said privacy preference is higher than its recovery rate, as expected with hypothesis 1a. This result shows support for using information diffusion models in the modelling of the spread of privacy preferences as an answer to RQ1. Furthermore, a positive correlation is found between the degree of an agent on the network and its influence on the spread of privacy preference, consistent with hypothesis 1b. This result also shows us that the most influential agents are the agents with the highest degrees in the network, which answers RQ2.

Rare privacy preference with high degree influencers Experiments in this setting are performed to investigate the ability of a few most influential users to spread rare privacy preferences. Rarity is defined as 0.2% of the network being infected with the privacy preference. In other words, 0.2% of the agents will be infected for the context tuples above. Furthermore, the top 1% of the agents, with regard to degree, are assigned the rare privacy preference. With this setting, it is hypothesized that:

- **2a** the influence of the top 1% of agents on the spread of the rare privacy preference will be significantly higher than in the baseline simulation setting.

After running the simulations, these were the results. The influence ratings of the best-connected agents rise when they are tasked with spreading a rare privacy preference in 76 out of 81 parameter settings, when opposing privacy preference

dynamics are equal to their unopposed counterparts. This hypothesis 2a holds in all parameter settings when opposing privacy preferences are assumed to be static. This result further supports the idea that the agents with the highest degree in the OSN are the most influential in the spread of privacy preferences, as an answer to RQ2.

Trust settings The following trust settings are simulated: (i) **Static trust** investigates how privacy preferences spread when agents trust each other to some degree and this level of trust remains static throughout an experiment. (ii) **Dynamic trust** investigates how privacy preferences spread when agents trust each other to some degree and this level of trust changes over time. (iii) **Dynamic trust with two levels of privacy violations** investigates how privacy preferences spread when agent model trust based on two types of privacy violations. Before running these simulations, simulations are run in which all trust values are 0 throughout the simulations. These simulations are used as sanity checks to ensure that without trust, no new infections are able to appear. For all three of these settings, the following claims are expected to hold:

- **Hypothesis 4a:** the epidemics will last shorter than in the baseline simulations, as the inclusion of trust negatively influences the infection rate.
- **Hypothesis 4b:** agents’ influence on the spread of privacy preferences will be diminished compared to the baseline simulations, since how much an agent is trusted by their neighbours now becomes a factor as well.
- **Hypothesis 4c:** epidemic peaks will be lower than in the baseline simulations.

An example of the state dynamics, when trust is introduced, can be seen in Figure 3b. It is noticeable that the decrease in the number of infected agents is quicker than baseline simulation where all rates are equal, Figure 3a. This is because now privacy violations are registered and they affect the general infection rate of the $\langle \textit{Night}, \textit{Work} \rangle$ privacy preference.

Simulation results show that when opposing privacy preference dynamics are equal to their unopposed counterparts, the introduction of trust modelling does not reduce the overall impact of privacy preference epidemics in all settings with regard to epidemics ending sooner and peaking lower. However, the influence of agents is reduced when trust modelling is introduced in the DIPP model. When opposing privacy preferences are assumed to be static, the introduction of trust modelling does reduce the impact of privacy preference epidemics. This result suggests that the introduction of trust modelling reduces the impact of privacy preference epidemics due to privacy violations, which answers RQ3.

5 Conclusion

Simulations using the DIPP model show that it provides a solid foundation to model the spread of privacy preferences. The well-connected agents are the most influential in the spread of privacy preferences in the DIPP model. The trust modelling allows agents to protect themselves against unwanted privacy

preferences. In the current model, an agent shares content and tries to recover from an infection at each step in the simulation. But this may not be realistic. In the state dynamics graphs, the number of agents that have recovered increase from step 0 in very similar fashion across different settings. A study of Facebook user sharing habits² found that the assumed uniformity in OSN use of all agents is not realistic. Our model can be used to configure the probability that agents take certain actions in further iterations of the DIPP model. Such extensions to the DIPP model can easily be made due to the modularity of our implementation. This study did not examine the influence of network structure on the spread of privacy preferences. It is clear that, although the network used is realistic, network structure can influence the results of our experiments. This is a point of attention for further research on this topic.

References

1. CANNARELLA, J., AND SPECHLER, J. A. Epidemiological modeling of online social network dynamics. *arXiv preprint arXiv:1401.4208* (2014).
2. DEGROOT, M. H. Reaching a consensus. *Journal of the American Statistical Association* 69, 345 (1974), 118–121.
3. DIGNUM, F. Autonomous agents with norms. *Artificial Intelligence and Law* 7, 1 (1999), 69–79.
4. DUPREE, J. L., DEVRIES, R., BERRY, D. M., AND LANK, E. Privacy personas: Clustering users via attitudes and behaviors toward security practices. In *Proceedings of the Conference on Human Factors in Computing Systems* (2016), pp. 5228–5239.
5. GUILLE, A., HACID, H., FAVRE, C., AND ZIGHED, D. A. Information diffusion in online social networks: A survey. *ACM Sigmod Record* 42, 2 (2013), 17–28.
6. JOSANG, A., AND ISMAIL, R. The beta reputation system. In *Proceedings of the 15th bled electronic commerce conference* (2002), vol. 5, pp. 2502–2511.
7. KÖKCIYAN, N., AND YOLUM, P. Priguard: A semantic approach to detect privacy violations in online social networks. *IEEE Transactions on Knowledge and Data Engineering* 28, 10 (2016), 2724–2737.
8. LAMPINEN, A., LEHTINEN, V., LEHMUSKALLIO, A., AND TAMMINEN, S. We’re in it together: interpersonal management of disclosure in social network services. In *Proceedings of the Conference on Human Factors in Computing Systems* (2011), pp. 3217–3226.
9. LI, M., WANG, X., GAO, K., AND ZHANG, S. A survey on information diffusion in online social networks: Models and methods. *Information* 8, 4 (2017), 118.
10. REECE, S., ROGERS, A., ROBERTS, S., AND JENNINGS, N. R. Rumours and reputation: Evaluating multi-dimensional trust within a decentralised reputation system. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems* (2007), pp. 1–8.
11. SAPIEZYNSKI, P., STOPCZYNSKI, A., LASSEN, D. D., AND LEHMANN, S. Interaction data from the copenhagen networks study. *Scientific Data* 6, 1 (2019), 1–10.
12. ULUSOY, O., AND YOLUM, P. Emergent privacy norms for collaborative systems. In *International Conference on Principles and Practice of Multi-Agent Systems* (2019), Springer, pp. 514–522.

² <https://www.frac.tl/work/marketing-research/facebook-user-sharing-habits-study/>